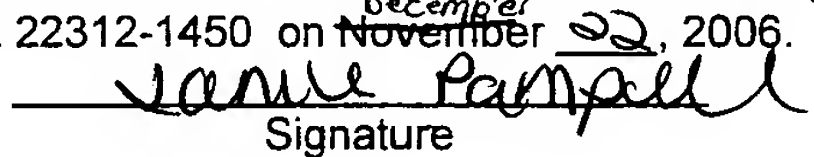
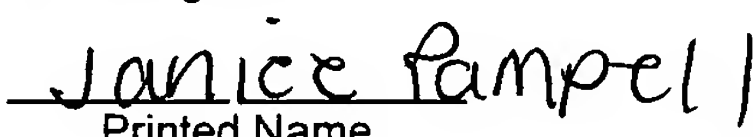
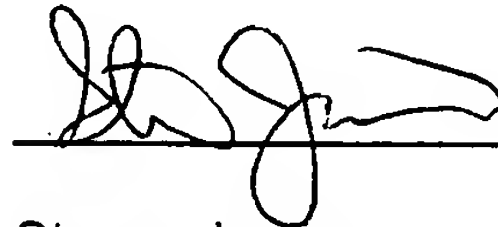
	
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE	
Declaration Under 37 C.F.R. 1.131	Atty. Docket No. CROSS1560
Applicants Steven A. Justiss	
Application Number 10/636,042	Date Filed 08/07/2003
Title System and Method for Maintaining and Reporting a Log of Multi-Threaded Backups	
Group Art Unit 2188	Examiner Hein, Gregory
Confirmation Number: 9977	

<p align="center"><u>Certificate of Mailing Under 37 C.F.R. §1.10</u></p> <p>I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22312-1450 on November ^{December} 22, 2006.</p> <p align="center">  Signature </p> <p align="center">  Printed Name </p>
--

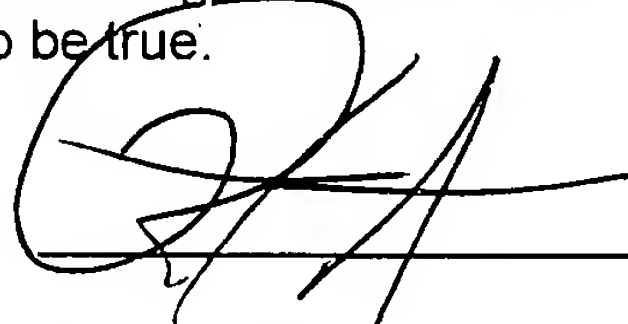
1. We, Robert Sims and Steve Justiss, employees of Crossroads Systems, Inc., are the original joint inventors of the invention described and claimed in the above-referenced patent application.
2. At least as early as July 12, 2002, we conceived the described the invention of the present application in an invention disclosure. A copy of the invention disclosure form evidencing conception at least as early as July 12, 2002 is attached hereto as Exhibit A.
3. In a letter dated February 21, 2003, Mark Berrier of Gray Cary, sent us a letter including a draft application describing the invention of the above-referenced application. A copy of the February 21, 2003 letter and draft application is attached hereto as Exhibit B. Support for the invention can be found at least at paragraphs 0010, 0011, 0012, 0042-0044, 0046 and 0053-0054 of the draft application.
4. The application was filed on August 7, 2003.

I, Steven Justiss, aver that all statements made of my own knowledge are true and all statements made on information and belief are believed to be true.

A handwritten signature in black ink, appearing to read 'Steven Justiss', written over a horizontal line.

Steven Justiss
November 3, 2006

I, Robert Sims, aver that all statements made of my own knowledge are true and all statements made on information and belief are believed to be true.

A handwritten signature in black ink, appearing to read 'Robert Sims', written over a horizontal line.

Robert Sims
November 8, 2006

EXHIBIT A



INVENTION DISCLOSURE FORM

PART ONE

I. Title

Method for Maintaining and Reporting a Log of Multi-Threaded Backups

II. Application

This invention applies to implementations of multi-threaded backup applications, including those that use the NCITS SCSI-2 Extended Copy command.

This invention applies to all microprocessor-based controllers that implement multi-threaded backups, including those that implement the NCITS SCSI-2 Extended Copy command

The invention applies to computer network backup technologies.

This invention applies to all forms of backup media (for example tape, or optical media).

III. Field of the Invention

The field of this invention is computer software.

IV. Background

Computer software data backup applications may initiate multiple backup threads that store data from different sources onto one backup medium (such as a tape, or optical disk). These backup applications may operate in a standalone configuration (storage devices directly attached to the host computer), or in a networked storage configuration (storage devices attached to a network, along with the host computer). The backup applications may transfer the data of backup threads directly to the storage media (such as in direct-attached, or in "lan-free" backups), or make use of "third party copy" backup strategies.

The NCITS T10 SPC-2 (SCSI Primary Commands-2) Extended Copy command provides a method for computer backup applications to delegate actual data movement to "third party devices" known as copy managers. These copy managers typically reside in mass storage related, microprocessor-based, storage network attached devices.

Copy managers move data from source devices to destination devices as designated by the backup application in "segment descriptors" which in part constitute the parameter list of the Extended Copy command. To enable

restoration of the data each block of real data is paired with "metadata" which contains identifying information about the real data, allowing its proper restoration from the backup medium.

When the backup medium is tape, the copy manager strives to keep a tape drive streaming (continuously moving the tape, writing data to the tape) in order to maximize performance. To keep the tape drive streaming the copy manager generally implements some form of disk data prefetch, or "read-ahead", so that the copy manager has data in buffers ready to build the next tape write command when an active tape write command completes.

The standard contemplates that a copy manager may handle some number of "concurrent" Extended Copy sessions depending on the size and number of system resources available. (The standard provides a method of reporting to the application the number of concurrent Extended Copy commands the copy manager can handle.)

An alternative means of keeping a tape drive streaming utilizes concurrent Extended Copy command-produced tape writes which are multi-threaded onto one tape (an invention described in a separate disclosure, namely "Method for Multi-threaded Extended Copy Backup to One Tape Drive"). Host applications may find it difficult to properly restore a backup medium written in this fashion.

This invention contemplates maintaining a log that records the source of write commands, and the order with which the backup medium is written. The source identification of the write command might consist of but is not limited to such identifiers as a protocol dependent Host ID, the XCopy-specification-defined List ID, a time stamp, and the size of the backup medium block written. The order with which the backup medium is written could be identified with these same Host ID and List ID numbers. Identifying the source and order with which a backup medium has been written could be utilized to properly restore archived data.

This invention further contemplates utilizing a system of vendor unique commands to perform policy based functions such as initiate a log (possibly including an application generated log identifier), retrieve a log (that is, send it from the copy manager to the host), copy of a log to a storage medium (such as append the log to the backup medium when the backup is complete), and clear a log (erase the log from copy manager memory). These log functions might be accomplished through purely vendor unique commands, or through a mix of vendor unique and Extended Copy vendor unique extensions.

V. Previous Solutions

Existing implementations do not provide a system for maintaining and/or reporting a log of write commands to storage media that were generated by one or several backup threads, such as concurrent Extended Copy commands.



VI. Summary of the Invention

This invention conceives maintaining a log of write commands generated by one or several backup streams that address one storage device. This allows the host application that generated the multi-threaded backup streams to decipher the data with which a storage medium has been written and thereby perform a proper restoration of the archived data.

VII. Advantages

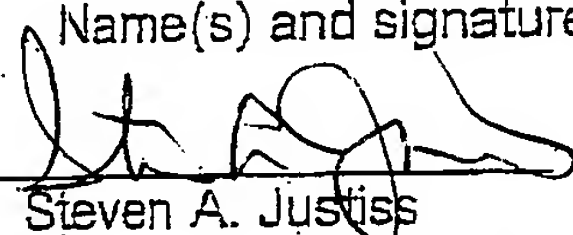
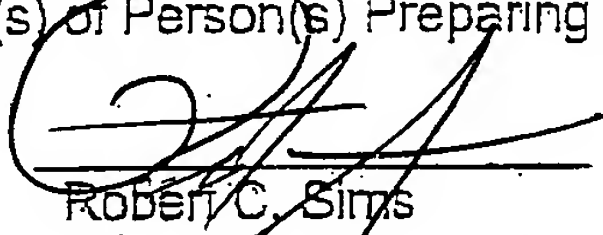
An application can utilize "multi-threading" of backup streams from multiple source disk drives onto a single tape drive as a mechanism to keep the tape drive streaming, and utilize the "backup log" to subsequently perform a restore. In addition the "backup log" might be used for diagnostic purposes, or for performance monitoring.

VIII. Disclosure Outside of Crossroads

This invention has not been disclosed outside Crossroads Systems.

IX. Inventorship

List inventors

	Name(s) and signature(s) of Person(s) Preparing this Form:		
Signature			
Name	Steven A. Justiss	Robert C. Sims	
Date	July 12, 2002	July 12, 2002	

1221 S. MoPac Expressway Suite 400
Austin, TX 78746-6875
www.graycary.com
O] 512-457-7016
F] 512-457-7001

February 21, 2003

VIA HAND DELIVERY

Mr. Steven A. Justiss
Mr. Robert Sims
Crossroads Systems, Inc.
8300 North MoPac Expressway
Austin, Texas 78759

Re: U.S. Patent Application Entitled "SYSTEM AND METHOD FOR MAINTAINING AND
REPORTING A LOG OF MULTI-THREADED BACKUPS"
Our Client/Docket No.: 103671.991560 (CROSS1560)

Dear Messrs. Justiss and Sims:

I enclose a copy of the above-identified application for patent, along with a redlined copy and new drawings. The application is now ready for execution and filing in the United States Patent and Trademark Office (PTO).

Please carefully review the application. If it accurately and adequately describes the invention, please execute the "Declaration and Power of Attorney" and "Assignment" documents, signing your name in blue ink, exactly as it is typewritten and dating each document. If any minor changes to the application are necessary, they should be made and such changes must be initialed and dated by you in the side margin closest to the changes, *before* signing the Declaration. If major changes are necessary, or if you have any questions, please call me. Also, the application must disclose the best mode of carrying out the invention; please let me know if it does not.

Please note that in executing the Declaration, you are acknowledging your duty to disclose material prior art to the PTO. Such prior art includes relevant patents and printed publications, information concerning public use of methods or apparatus related to your invention, and information on public use or sales of your own invention (or related methods or apparatus) made more than a year ago. Your failure to disclose such prior art may invalidate any patent issuing on the application.

Once these documents have been executed, please return the application, executed Declaration and Power of Attorney and Assignment to me in the enclosed self-addressed stamped envelope in order that we may file the application with the PTO as soon as possible.

SILICON VALLEY SAN DIEGO SAN DIEGO/GOLDEN TRIANGLE SAN FRANCISCO AUSTIN SEATTLE SACRAMENTO LA JOLLA

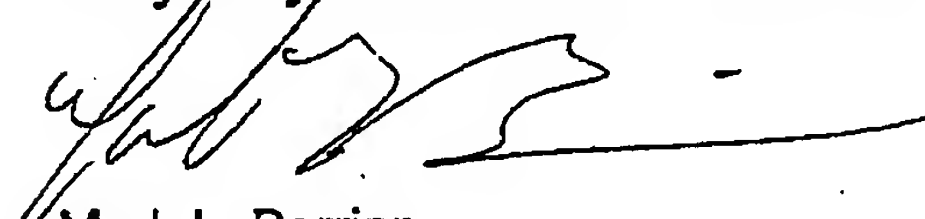
Gray Cary\AU\4098586.1
103671-991560

Messrs. Justiss and Sims
February 21, 2003
Page 2

Should there be any questions concerning this matter, please feel free to contact me at (512) 457-7016.

Sincerely,

Gray Cary Ware & Freidenrich LLP



Mark L. Berrier

MLB:klc
Enclosures

CORPORATE, FINANCIAL AND INDUSTRIAL DELIVERY
1000 EAST 7TH ST., SUITE 201
AUSTIN, TX 78702
(512) 320-5448

GRAY CARY WARE & FREIDENRICH, LLP
1221 S. MOPAC STE. 400
AUSTIN, TEXAS 78746 (512) 457-7000

F R O M		TO		DATE	
PJ 1:15pm		2.24.03		ORIGINAL	
Steven Justiss/Rob Sims					
Crossroads, 8300 N. Mopac					
Austin, TX 78759					
REF. NO.		CUST. NO.			
103671.991560		103671.991560			
BILL TO		BANKRUPTCY		COUNTY	
DT/DT		FEDERAL		PROBATE	
DIRECT		HOT SHOT		REGULAR	
ASAP		COURTHOUSE RUN		DISTRICT	
NO PCS		DESCRIPTION		WEIGHT	
		Application & formal			
DRIVER		/		DRIVER:	
		1 RECEIVED IN GOOD ORDER:		TIME	
				DELIVERY CHARGE	
				RETURN	
				WAITING TIME	
				OTHER	
				TOTAL	

Local or damage claims may not exceed \$50 unless a greater value is declared. Maximum declared value \$2,000.
CLAIMS MUST BE MADE WITHIN 48 HOURS. NOT RESPONSIBLE FOR CONCEALED DAMAGE.

ASSIGNMENT

This Assignment is made by **Steven A. Justiss, of Lakeway, Texas and Robert Sims of Round Rock, Texas** ("Assignors").

WHEREAS, Assignors have invented a new and useful invention entitled **SYSTEM AND METHOD FOR MAINTAINING AND REPORTING A LOG OF MULTI-THREADED BACKUPS**, for which an application for United States Patent is made, said application being submitted concurrently herewith; and

WHEREAS, Assignors believe themselves to be the original inventors of the invention including any and all improvements disclosed in said application ("Invention"); and

WHEREAS, the parties desire to have a recordable instrument assigning the entire right, title and interest in and to said Invention, said application and any patents, invention registrations or other forms of protection ("Patents") that may be granted for said inventions in the United States and throughout the world;

NOW, THEREFORE, in accordance with the obligations to assign the Invention and for other good and valuable consideration, the receipt and sufficiency of which are hereby acknowledged, Assignors hereby sell, assign, and transfer to **Crossroads Systems, Inc.** having a principal place of business at **8300 North MoPac Expressway, Austin, Texas 78759** (hereinafter referred to as "Assignee"), the entire right, title, and interest in and to said Invention, said application and any Patents that may be granted for said Invention in the United States and throughout the world, including the right to file foreign applications directly in the name of the Assignee and to claim for any such foreign applications any priority rights to which such applications are entitled under international conventions, treaties, or otherwise.

Assignors agree that, upon request and without further compensation, but at no expense to Assignors, he/she and/or their legal representatives and assigns will do all lawful acts, including the execution of papers and the giving of testimony, that may be necessary or desirable for obtaining, sustaining, reissuing, or enforcing the Patents in the United States and throughout the world for said Invention, and for perfecting, recording, or maintaining the title of Assignee, its successors and assigns, to said Invention, said application, and any Patents granted for said Invention in the United States and throughout the world.

Assignors represent and warrant that they have not granted and will not grant to others any rights inconsistent with the rights granted herein.

Assignors authorize and request the Assistant Commissioner for Patents of the United States and of all foreign countries to issue any Patents granted for said Invention,

whether on said application or on any subsequently filed divisional, continuation, continuation-in-part, reissue or other application, to Assignee, its successors and assigns, as the assignee of the entire interest in said invention.

IN WITNESS WHEREOF, Assignors have executed this Assignment on the dates provided below.

Name of First Inventor:

Steven A. Justiss

Signature:



Date:

Nov. 3, 2006

Citizenship

United States of America

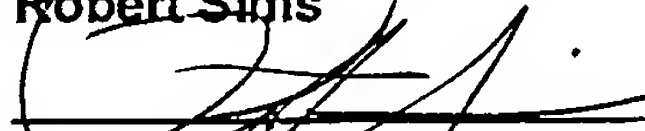
Residence Address:

603 Brooks Hollow Road
Lakeway, TX 78734

Name of Inventor:

Robert Sims

Signature:



Date:

Nov. 8, 2006

Citizenship

8609 Sea Ash Circle

Residence Address:

Round Rock, Texas 75681

**DECLARATION FOR
UTILITY OR DESIGN
PATENT APPLICATION
(37 CFR 1.63)**

Attorney Docket No.

CROSS1560

First Named Inventor

Justiss et al.

COMPLETE IF KNOWN

Application Number

Unknown

Filing Date

Unknown

Group Art Unit

Unknown

Examiner Name

Unknown

☒ Declaration Submitted
with Initial Filing☐ Declaration Submitted after
Initial Filing**As a below named inventor, I hereby declare that:**

My residence, post office address, and citizenship are as stated below to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

SYSTEM AND METHOD FOR MAINTAINING AND REPORTING A LOG OF MULTI-THREADED BACKUPS

(Title of Invention)

the specification of which was filed on (MM/DD/YYYY)

as United States Application Number or PCT International
Application Number

and was amended on (MM/DD/YYYY) (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment specifically referred to above.

I acknowledge the duty to disclose to the U.S. Patent and Trademark Office all information known to me which is material to the patentability as defined in 37 CFR 1.56, including for continuation-in-part applications, material information which became available between the filing date of the prior application and the national or PCT international filing date of the continuation-in-part.

I hereby claim foreign priority benefits under 35 U.S.C. 119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate, or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or of any PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application Number(s)	Country	Foreign Filing Date (MM/DD/YYYY)	Priority Not Claimed	Certified Copy Attached? YES NO

Additional foreign application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto:

I hereby claim the benefit under 35 U.S.C. 119(e) of any United States provisional application(s) listed below:

Application Number(s)	Filing Date (MM/DD/YYYY)	
		<input type="checkbox"/> Additional provisional application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto

DECLARATION - Utility or Design Patent Application

I hereby claim the benefit under 35 U.S.C. 120 of any United States Application(s), or 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. 112, I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application.

U.S. Parent Application or PCT Parent Number	Parent Filing Date (MM/DD/YYYY)	Parent Patent Number (If applicable)

Additional U.S. or PCT international application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto.

As a named inventor, I hereby appoint the registered practitioner(s) assigned to **Customer No. 25094** to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Name of Sole/First Inventor:

Given Name (first and middle [if any])

Family Name or Surname

Steven A.

Justiss

Inventor's
Signature

Date

Residence: City

Lakeway

State

TX

Country

United States

Citizenship

United States

Residence Address

603 Brooks Hollow Road, Lakeway, TX 78734

Post Office Address

same

Name of Additional Inventor:

Given Name (first and middle [if any])

Family Name or Surname

Robert

Sims

Inventor's
Signature

Date

Residence: City

Round Rock

State

TX

Country

United States

Citizenship

Residence Address

8609 Sea Ash Circle, Round Rock, TX 75681

Post Office
Address

same

Name of Additional Inventor:

Given Name (first and middle [if any])

Family Name or Surname

Inventor's
Signature

Date

Residence: City

State

Country

Citizenship

Residence Address

Post Office Address

Direct all correspondence to Customer No. 25094:

Name

Mark L. Berrier

Gray Cary Ware & Freidenrich LLP

Address

1221 So. MoPac Expressway, Suite 400

City

Austin

State

TX

Zip

78746

Country

U.S.A.

Telephone

(512) 457-7016

Fax

(512) 457-7001

- 1 -

**SYSTEM AND METHOD FOR MAINTAINING AND
REPORTING A LOG OF MULTI-THREADED BACKUPS**

FIELD OF THE INVENTION

[0001] The invention relates generally to storage and retrieval of data and more particularly to improved systems and methods for retrieving data from a sequential storage device, where the data was contained in one of multiple threads that were stored on the device in an intermingled fashion.

RELATED ART

- [0002] The NCITS T10 SPC-2 (SCSI Primary Commands-2) Extended Copy command provides a method for computer backup applications to delegate actual data movement to "third party devices" known as copy managers. These copy managers typically reside in mass storage related, microprocessor-based, storage network attached devices.
- [0003] Copy managers move data from source devices to destination devices as designated by the backup application in "segment descriptors" which in part constitute the parameter list of the Extended Copy command. To enable restoration of the data each block of real data is paired with "metadata" which contains identifying information about the real data, allowing its proper restoration from the backup medium.
- [0004] The standard contemplates that a copy manager may handle some number of "concurrent" Extended Copy sessions depending on the size and number of system resources available. (The standard provides a method of reporting to the application the number of concurrent Extended Copy commands the copy manager can handle.)
- [0005] Computer software data backup applications may initiate multiple backup threads that store data from different sources onto one backup medium (such as a tape, or optical disk). These backup applications may operate in a standalone configuration (storage devices directly attached to the host computer), or in a networked storage configuration (storage devices attached to a network, along with the host computer). The backup applications may transfer the data of backup threads directly to the storage media (such as in direct-attached, or in "lan-free" backups), or make use of "third party copy" backup strategies.
- [0006] When the backup medium is tape, the copy manager strives to keep a tape drive streaming (continuously moving the tape, writing data to the tape) in order to maximize performance. To keep the tape drive streaming the copy manager generally implements some form of disk data prefetch, or "read-ahead", so that the copy

- 3 -

manager has data in buffers ready to build the next tape write command when an active tape write command completes.

[0007] An alternative means of keeping a tape drive streaming utilizes concurrent Extended Copy command-produced tape writes which are multi-threaded onto one tape (an invention described in a separate disclosure, namely "Method for Multi-threaded Extended Copy Backup to One Tape Drive"). Host applications may find it difficult to properly restore a backup medium written in this fashion.

SUMMARY

- [0008] One or more of the problems outlined above may be solved by the various embodiments of the invention. Broadly speaking, the invention comprises systems and methods for performing multi-threaded backups and restores.
- [0009] In one embodiment, a log is maintained to record the source of write commands, and the order in which blocks of data are written to the backup medium. The source identification of the write command may consist of such identifiers as a protocol dependent Host ID, the extended-copy-specification-defined List ID, a time stamp, and the size of the backup medium block written. The order in which the data is written to the backup medium can be identified with these same Host ID and List ID numbers. In other words, as blocks of data from each thread are written, the thread and the number of blocks is identified in the log. Then, when it is desired to restore data corresponding to one of the threads, the desired blocks of data can be identified in the log, and the preceding blocks stored on the backup medium can be skipped.
- [0010] One embodiment of the invention is a method comprising generating a write log, wherein the write log identifies a sequence in which blocks of data corresponding to multiple write threads are stored on a sequential device, reading the log, identifying at least a portion of the blocks of data corresponding to one of the write threads and indexing to the location of the identified portion of the blocks of data in the sequence of blocks of data stored onto the sequential device according to the write log. Another embodiment comprises a method similar to that described above, but wherein the write log was previously generated, and the method comprises reading the log, identifying the desired data in the log, indexing to the location of the data on the sequential device as indicated in the log and retrieving the data from the sequential device.
- [0011] One embodiment of the invention comprises a system for managing blocks of data on a sequential storage device, wherein blocks of data corresponding to multiple threads are stored on the sequential storage device in an intermingled fashion, comprising a sequential storage device configured to store intermingled blocks of data

corresponding to multiple threads and a copy manager coupled to the sequential storage device and configured to manage the retrieval of copying of desired blocks of data from the sequential storage device. The system may also include a memory coupled to the copy manager and configured to store a sequence in which blocks of data corresponding to multiple threads are stored on the sequential storage device. The copy manager is configured to identify the position of the desired blocks of data in the sequence stored in the memory, to advance to a corresponding storage location on the sequential storage device without reading each of the preceding stored blocks of data, and to retrieve the desired blocks of data from the sequential storage device.

[0012] Another embodiment of the invention comprises a software application. The software application is embodied in a storage medium readable by a computer or other data processor, such as a floppy disk, CD-ROM, DVD-ROM, RAM, ROM, or the like. The storage medium contains instructions which are configured to cause a data processor such as a router or other SAN component to execute a method which is generally as described above. It should be noted that the storage medium may comprise a RAM or other memory which forms part of a data processor. The data processor would thereby be enabled to perform a method in accordance with the present disclosure.

[0013] Numerous additional embodiments are also possible.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0014] Other objects and advantages of the invention may become apparent upon reading the following detailed description and upon reference to the accompanying drawings.
- [0015] FIGURE 1 is a diagram illustrating the interconnection of the network components in one embodiment.
- [0016] FIGURE 2 is a diagram illustrating the general structure of an extended copy command in one embodiment.
- [0017] FIGURES 3A and 3B are a pair of diagrams illustrating the flow of extended copy commands and corresponding data flow in one embodiment.
- [0018] FIGURE 4 is a diagram illustrating the multiplexing of the data streams resulting from different threads of execution of different extended copy commands is shown.
- [0019] FIGURE 5 is a diagram illustrating the manner in which blocks of data from different threads are multiplexed to form a single stream of blocks to be written to a sequential storage device.
- [0020] While the invention is subject to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and the accompanying detailed description. It should be understood, however, that the drawings and detailed description are not intended to limit the invention to the particular embodiment which is described. This disclosure is instead intended to cover all modifications, equivalents and alternatives falling within the scope of the present invention as defined by the appended claims.

DETAILED DESCRIPTION

- [0021] A preferred embodiment of the invention is described below. It should be noted that this and any other embodiments described below are exemplary and are intended to be illustrative of the invention rather than limiting.
- [0022] Broadly speaking, the invention comprises systems and methods for performing storage and/or retrieval of data blocks corresponding to multiple threads, where the blocks of data corresponding to different threads are intermingled on a sequential storage device.
- [0023] One embodiment of the invention is implemented in the switching fabric of a storage area network (SAN). Referring to FIGURE 1, a block diagram illustrating the interconnection of the network components in one embodiment is shown. In this embodiment, host 12 is coupled to a switch fabric 14. Multiple devices such as disk drives 17, 18 and 19, as well as sequential storage device 16 are also coupled to switch fabric 14.
- [0024] Host 12 may be any of a variety of devices, such as a Solaris box, a Windows 2000 server, or the like. Devices 17-19 may likewise comprise several types of devices from which data may need to be copied. For example, in one embodiment, devices 17-19 may all be hard disk drives that need to be backed up. Sequential storage device 16 may also be one of several different types of storage devices. Typically, sequential storage device 16 is a tape drive. In an alternative embodiment, it might be an optical drive.
- [0025] It should be noted that switch fabric 14 may have many different configurations. For example, it may comprise a router or various other types of SAN attached appliances. Switch fabric 14 may be coupled to host 12, devices 17-19 and sequential storage device 16 via various types of interconnects. For instance, they may comprise Fibre Channel, SCSI, iSCSI, InfiniBand, or any other type of interconnect that allows transport of INCITS T10 SCSI extended copy commands.

- [0026] Host 12 is capable of accessing the other components of the network via switch fabric 14. Particularly, host 12 is capable of accessing a copy manager application resident in one of the components of switch fabric 14 to delegate to it the management of copy tasks involving other network components. For example, host 12 might delegate the task of backing up data from devices 17-19 to sequential storage device 16. This is accomplished through the use of extended copy commands issued by host 12 to the copy manager. The copy manager executes these extended copy commands, reading data from devices 17-19 and writing (copying) the data to sequential storage device 16.
- [0027] The use of extended copy commands allows host 12 to use its own processing power on tasks other than the mere movement of data between network components. For example, it is possible for host 12 to back up a hard disk drive to a tape drive (a potentially very lengthy process) by issuing one or more corresponding extended copy commands to the copy manager. The copy manager can then copy the backup data from the hard disk drive to the tape drive without the intervention of host 12.
- [0028] Referring to FIGURE 2, a diagram illustrating the general structure of an extended copy command in one embodiment is shown. As shown in the figure, the extended copy command 30 has an opcode 32. In this instance, the opcode is a hexadecimal "83". The command format includes a link or pointer list 34 which is a count of the size of a list 36 in the data. List 36 has a header 41, a set of target descriptors 42, a set of segment descriptors 43 and in-line data 44.
- [0029] The target descriptors 42 describe the target devices which will be involved in the extended copy task. The target devices normally include the source device and the destination device (i.e., the device from which data will be read and the device to which the data will be written). There may, however, be additional target descriptors, and the preferred embodiment makes provision for up to 64 target descriptors. Parameters in the target descriptor list may include such things as the address, name, size of data blocks, fixed-/variable-block mode, etc. for the target device.

- [0030] The segment descriptors 43 describe the types of operations that will be performed, as well as the amount of data that will be transferred. For example, a segment descriptor may indicate that a block of data will be read from a hard disk drive and written to a tape drive (both of which are referenced by the target descriptors), that is, a backup operation. Alternatively, the segment descriptors may describe backup operations, restore operations, block-to-block operations, etc. It should be noted that inline data (which may also be referred to as metadata) may or may not be present. Typically, for operations such as the backup of a disk to tape, the inline data is present. Metadata may also be present as SPC-2-defined "embedded data". A preferred embodiment is an SPC-2 implementation which supports 8448 segment descriptors. Each of the segment descriptors can move up to 32 MB of data. Thus, a very large amount of data may potentially be moved through a single extended copy command.
- [0031] In the preferred embodiment, the data is not read or written in 32 MB chunks. The data is instead read from a disk in chunks of 256 kB and then written to a tape in chunks of 64 kB. Internal buffers used in the preferred embodiment are 16 kB each, so each read command to the disk requires 16 buffers and each write command to the tape requires four buffers. This information is used to determine whether sufficient resources are available to activate additional extended copy commands.
- [0032] In one embodiment, the network component in which the copy manager resides employs buffers to store data that is read from the source device prior to writing it to the destination device. In this embodiment, the buffers in which the data is stored have a low-water mark associated with them. After one or more read commands are issued and the corresponding data fills the buffers, write commands are issued to transfer data from the internal buffers to the destination device. These write operations continue until the number of buffers which contain data (which has not yet been transferred to the destination device) falls below the predetermined low-water mark. When the level of data in the buffers falls below the low-water mark, additional read commands are issued to obtain more data and store it in the buffers.

- [0033] Referring to FIGURES 3A and 3B, a pair of diagrams illustrating the flow of extended copy commands and corresponding data flow are shown. FIGURES 3A and 3B depict the flow of commands and data through the system during a backup operation that is performed using extended copy commands. In this instance, data from all three devices 17-19 is backed up to sequential storage device 16.
- [0034] Data flow A (shown in FIGURE 3A) represents the issuance of extended copy commands from the host 12 to a copy manager 13. The copy manager can be implemented in a router, or in some other component of the SAN. Three different extended copy commands are issued in the depicted situation – one for each of devices 17-19 to be backed up. Each of the extended copy commands instructs the copy manager to back up data from one of the devices to the sequential storage device 16. When copy manager 13 receives each extended copy command, it constructs a corresponding series of commands that will be issued to the appropriate device to read data from it. The copy manager also constructs a series of commands that will be issued to sequential storage device 16 to write the devices' data to it. The read commands issued by the copy manager to each device are shown as data flows B₁, B₂ and B₃ in FIGURE 3A. The data responsive to these read commands is shown as data flows C₁, C₂ and C₃ in FIGURE 3B. Finally, the write commands issued by the router to the tape drive are shown as dataflow D (see FIGURE 3B).
- [0035] When the extended copy command is executed, the copy manager identifies the target devices and goes through the segment descriptors sequentially (in the order received). For each of the segment descriptors, the copy manager builds corresponding read commands to be issued to the source device and write commands to be issued to the destination device. When these read and write commands are issued to the respective devices, they serve to transfer data from the source device to the destination device.
- [0036] If there is only one extended copy command that is being executed at a time, the copy manager simply receives the data from the source device and transfers the data to the destination device. If the data is being stored on a sequential storage device, the data is recorded on the storage device in the same sequence that it is received. In one

embodiment, the data is received by the sequential storage device in blocks, so consecutive blocks are stored in consecutive (adjacent) locations in the device's storage medium. Then, when it is necessary to retrieve the data from the sequential storage device, a series of consecutive blocks corresponding to the desired data is read from the device.

[0037] It should be noted that the term "block," as used herein should be construed to include blocks of data that have any appropriate size and/or formatting. In fact, it is not necessary that all of the blocks have the same formatting, as long as the blocks can be delineated.

[0038] The situation becomes more complicated when there are multiple extended copy commands that are being executed at the same time. For example, multiple copy commands may be issued to the copy manager to transfer data from multiple devices to the same sequential storage device. In this case, the transfers of data from different devices to the copy manager overlap. The copy manager then transfers the received data in an intermingled fashion to the sequential storage device.

[0039] This situation should be distinguished from concurrent extended copy commands as referenced in the NCITS T10 SPC-2 Extended Copy Command specification. In this specification, "concurrent" extended copy commands are commands that have been issued and are pending at the same time, but they are not executed simultaneously. As commonly implemented, SPC-2 extended copy commands read from different source disks and write to different destination tape drives in an interleaved fashion.

[0040] The present systems and methods deal with the concurrent execution situation. As noted above, the data received from different devices in this situation is intermingled. More specifically, the data is multiplexed. Referring to FIGURE 4, a diagram illustrating the multiplexing of the data streams resulting from different threads of execution of different extended copy commands is shown.

[0041] FIGURE 4 depicts three streams of data, 51-53. Each stream corresponds to a particular extended copy command which is concurrently executing in its own thread.

- 12 -

As each thread executes, blocks of data are read from the corresponding device and transferred to the copy manager 13. The individual blocks of data are depicted in FIGURE 5. The blocks corresponding to the different threads may arrive at the copy manager at different rates, and they are not synchronized with the blocks of the other threads.

[0042] Referring to FIGURE 5, as the blocks of data 61-63 from the different threads 51-53 are received by the copy manager, they are multiplexed to form a single stream of blocks 65 and queued to be written to the sequential storage device. Because the blocks are randomly received, they are randomly intermingled. In other words, while all of the blocks corresponding to a given thread will be in order, the order in which the blocks alternate between the different threads is indeterminate. The number of consecutive blocks in the queue from a single one of the threads between blocks from other threads is also unknown.

[0043] Referring again to FIGURE 4, as the blocks of data are written from the queue to sequential storage device 16, however, the order of the blocks is tracked. In one embodiment, this is done by logging the write commands to sequential storage device 16 (i.e., by writing corresponding entries to log 20). Each entry in the log 20 includes an identifier of the thread to which the corresponding blocks belong, as well as the number of blocks from that thread. Each entry in log 20 may also contain information that may be useful for performance or diagnostic analyses. For example, each entry may include a timestamp that identifies the time at which the data was received or written. This information could then be used to determine the rate at which the data transfer occurred.

[0044] It should be noted that that, in other embodiments, the order in which the blocks are stored on the sequential storage device may be tracked in a different manner. For example, rather than recording the number of consecutive blocks from a particular thread, each block may be logged separately. In other words, consecutive blocks from a particular thread may have consecutive entries in the log, each having the same

- 13 -

thread identifier. In another embodiment, an identifier other than the thread identifier may be recorded with the entries in the log.

[0045] The primary purpose of logging the writes to the sequential storage device is to provide a mechanism for efficiently retrieving the recorded data. Although it is possible to read each of the individual recorded blocks of data to determine their respective origins, it is much more efficient to be able to simply index into the series of recorded blocks to find the particular blocks that are desired. The write log provides the information necessary to do this.

[0046] When it is desired to retrieve a portion of the data recorded to the sequential storage device from multiple threads, it is first necessary to identify which thread and which block (or blocks) of that thread contains the desired data. That block can then be identified in the write log. (As noted above, all of the blocks of a given thread are recorded in order.) Once the block is identified in the write log, it is only necessary to index into the recorded sequence of blocks to locate the desired block. (For example, the number of blocks that precede the desired block in the recorded series of blocks can be determined and this number of blocks can be skipped from the beginning of the series.) The desired block can then be read.

[0047] A restore operation using the write log can itself be implemented using extended copy commands. The host could read the write log and then construct an extended copy command having the following parameter list.

<u>Segment descriptor</u>	<u>Function</u>
1	Space to block N
2	Read block N
3	Space to block N+1
4	Read block N+1
5	Space to block N+J
6	Read block N+J
...	

- [0048] While it is possible to read the desired block of data from the sequential storage device without using the write log, this typically is not very efficient. For example, the *n*th block of stream *Y* can be retrieved by reading all of the preceding blocks, rather than simply skipping over them. If there are only relatively few blocks of data corresponding to stream *Y*, most of the resources used to retrieve the desired block of data will be wasted reading unwanted blocks. Clearly, this is inefficient.
- [0049] It is also possible to read metadata in each of the blocks to identify them. When the desired block is found, it is identified and read from the sequential storage device. While this may be more efficient than reading the preceding data blocks in their entirety, it is still inefficient in that the identifying metadata for each of the preceding blocks must be read. When the write log is used to identify the location of the desired block, it is not necessary to read the metadata in order to identify the desired block.
- [0050] Some of the embodiments of the invention contemplate utilizing a system of vendor unique commands to perform policy based functions. These may include functions such as initiating a log (possibly including an application-generated log identifier), retrieving a log (that is, sending it from the copy manager to the host), copying of a log to a storage medium (such as appending the log to the backup medium when the backup is complete), and clearing a log (erasing the log from copy manager memory). These log functions might be accomplished through purely vendor unique commands, or through a mix of vendor unique and Extended Copy vendor unique extensions.
- [0051] In addition to the embodiments described above, there are numerous alternative embodiments. For example, while the write log described above records a thread/device identifier and a number of consecutive blocks corresponding to that thread, it is possible to record the identifier for each block recorded. It would still be possible to use this log to determine how many blocks to skip from the beginning of the recorded sequence in order to arrive at the location of the desired block.

[0052] In another alternative embodiment, the sequence of the recorded data blocks may be tracked by a mechanism other than the write log described above. For instance, a log may be generated from the blocks of data in queue to be written to the sequential storage device, since they are written to the sequential storage device in the order they are retrieved from the queue. In yet another embodiment, the order of the data blocks may be recorded as the blocks are received, assuming that the blocks are written to the sequential storage device in the order they are received. This last alternative is not preferred, however, as two or more blocks may be received at the same time.

[0053] Other alternative embodiments may comprise methods employed by the systems described above. For example, One embodiment of the invention is a method comprising generating a write log, wherein the write log identifies a sequence in which blocks of data corresponding to multiple write threads are stored on a sequential device, reading the log, identifying at least a portion of the blocks of data corresponding to one of the write threads and indexing to the location of the identified portion of the blocks of data in the sequence of blocks of data stored onto the sequential device according to the write log. Another embodiment comprises a method similar to that described above, but wherein the write log was previously generated, and the method comprises reading the log, identifying the desired data in the log, indexing to the location of the data on the sequential device as indicated in the log and retrieving the data from the sequential device.

[0054] It should be noted that the methodologies disclosed herein may be implemented in various combinations of software (including firmware) and hardware. The present application is therefore intended to cover software applications that include instructions for causing a computer or other data processor to perform the methods disclosed herein. These software

- 16 -

applications may be embodied in any medium readable by such a computer or data processor, including floppy disks, CD-ROMs, DVD-ROMs, RAM, ROM, and the like. Likewise, a computer or data processor which is configured to execute such software applications, or which is otherwise programmed to perform the methods disclosed herein is intended to be covered by the present application.

[0055] The benefits and advantages which may be provided by the present invention have been described above with regard to specific embodiments. These benefits and advantages, and any elements or limitations that may cause them to occur or to become more pronounced are not to be construed as a critical, required, or essential features of any or all of the claims. As used herein, the terms 'comprises,' 'comprising,' or any other variations thereof, are intended to be interpreted as non-exclusively including the elements or limitations which follow those terms. Accordingly, a process, method, article, or apparatus that comprises a list of elements does not include only those elements but may include other elements not expressly listed or inherent to the claimed process, method, article, or apparatus.

[0056] While the present invention has been described with reference to particular embodiments, it should be understood that the embodiments are illustrative and that the scope of the invention is not limited to these embodiments. Many variations, modifications, additions and improvements to the embodiments described above are possible. It is contemplated that these variations, modifications, additions and improvements fall within the scope of the invention as detailed within the following claims.

WHAT IS CLAIMED IS:

1. A method for retrieving data from a sequential storage device on which blocks of data corresponding to multiple threads are stored in an intermingled fashion, comprising:
 reading a log, wherein the log identifies a sequence in which blocks of data
 corresponding to multiple threads are stored on a sequential storage device;
 identifying at least a portion of the blocks of data corresponding to one of the threads;
 and
 indexing to the location of the identified portion of the blocks of data in the sequence of
 blocks of data stored on the sequential device according to the log.
2. The method of claim 1, wherein indexing to the location of the identified portion of the blocks of data in the sequence comprises counting a first number of blocks preceding the identified portion of the blocks in the log and advancing the first number of blocks on the sequential device.
3. The method of claim 2, further comprising retrieving the identified portion of the blocks of data from the sequential storage device.
4. The method of claim 1, wherein the log includes indications of file marks in the stored blocks of data and wherein indexing to the location of the identified portion of the blocks of data in the sequence comprises counting a first number of file marks preceding the identified portion of the blocks in the log and advancing the first number of file marks on the sequential device.
5. The method of claim 4, further comprising retrieving the identified portion of the blocks of data from the sequential storage device.
6. The method of claim 1, further comprising storing the blocks of data on the sequential storage device and writing the log prior to reading the log.

- 18 -

7. The method of claim 1, wherein each entry in the log identifies a corresponding thread and a number of blocks stored consecutively on the sequential storage device.
8. The method of claim 7, wherein each thread is identified by a corresponding device identifier.
9. The method of claim 1, wherein the log is stored on the sequential storage device and the log is read from the sequential storage device.
10. The method of claim 1, wherein the log is stored on a storage medium separate from the sequential storage device and the log is read from the separate storage medium.

- 19 -

11. A method for managing storage of blocks of data on a sequential storage device, wherein blocks of data corresponding to multiple threads are stored on the sequential storage device in an intermingled fashion, comprising:

storing a sequence of blocks of data on a sequential storage device, wherein the blocks of data correspond to multiple write threads and wherein blocks corresponding to different write threads are intermingled on the sequential storage device;

recording the order in which the blocks of data are stored in a log; and
storing the log.

12. The method of claim 11, wherein recording the order in which the blocks of data are stored comprises recording entries corresponding to write commands in the log.

13. The method of claim 11, further comprising storing the log on the sequential storage device.

14. The method of claim 11, further comprising storing the log on a storage medium which is separate from the sequential storage device.

15. The method of claim 11, further comprising identifying at least a portion of the blocks of data corresponding to one of the threads, identifying the position of entries corresponding to the identified portion of the blocks of data in the log, and indexing to the location of the identified portion of the blocks of data in the sequence of blocks of data stored on the sequential device based upon the identified portion of the blocks of data in the log.

- 20 -

16. A system for managing blocks of data on a sequential storage device, wherein blocks of data corresponding to multiple threads are stored on the sequential storage device in an intermingled fashion, comprising:

- a sequential storage device configured to store intermingled blocks of data corresponding to multiple threads;

- a copy manager coupled to the sequential storage device and configured to manage the retrieval of copying of desired blocks of data from the sequential storage device; and

- a memory coupled to the copy manager and configured to store a sequence in which blocks of data corresponding to multiple threads are stored on the sequential storage device;

wherein the copy manager is configured to identify the position of the desired blocks of data in the sequence stored in the memory, to advance to a corresponding storage location on the sequential storage device without reading each of the preceding stored blocks of data, and to retrieve the desired blocks of data from the sequential storage device.

17. The system of claim 16, wherein the copy manager is further configured to store the sequence of the stored data blocks in the memory.

18. The system of claim 16, wherein the copy manager is configured to copy data to the sequential storage device according to a plurality of extended copy commands.

19. The system of claim 16, further comprising one or more hosts coupled to the copy manager, wherein the copy manager is configured to store the blocks of data on the sequential storage device according to extended copy commands issued by the one or more hosts.

20. The system of claim 19, further comprising a plurality of data sources, wherein the copy manager is configured to copy data from each of the plurality of data sources in a plurality of corresponding threads.

- 21 -

21. The system of claim 16, wherein the copy manager is implemented in a switch fabric.
22. The system of claim 16, wherein the copy manager is implemented in a network attached device.

- 22 -

23. A software product comprising one or more instructions embodied in a medium readable by a data processor, wherein the instructions are configured to cause the data processor to execute the method comprising:

- reading a log, wherein the log identifies a sequence in which blocks of data corresponding to multiple threads are stored on a sequential storage device;
- identifying at least a portion of the blocks of data corresponding to one of the threads;
- and
- indexing to the location of the identified portion of the blocks of data in the sequence of blocks of data stored on the sequential device according to the log.

**SYSTEM AND METHOD FOR MAINTAINING AND
REPORTING A LOG OF MULTI-THREADED BACKUPS**

ABSTRACT OF THE DISCLOSURE

[0057] Systems and methods for performing multi-threaded backups and restores. In one embodiment, a log is maintained to record the source of write commands, and the order in which blocks of data are written to a sequential storage device. The source identification of the write command may consist of such identifiers as a protocol dependent Host ID, the extended-copy-specification-defined List ID, a time stamp, and the size of the backup medium block written. The order in which the data is written to the backup medium can be identified with these same Host ID and List ID numbers. When it is desired to restore data corresponding to one of the threads, the desired blocks of data can be identified in the log, and the preceding blocks stored on the backup medium can be skipped.




Fig. 1




Fig. 2




Fig. 3A




Fig. 3B




Fig. 4




Fig. 5